Utilitarian Distributed Constraint Optimization Problems

Julien Savaux, Julien Vion, Sylvain Piechowiak, René Mandiau[1] and
Toshihiro Matsui[2] and **Katsutoshi Hirayama**[3] and **Makoto Yokoo** [4] and **Shakre Elmane, Marius Silaghi** [5]

arXiv:1604.06787v1 [cs.AI] 22 Apr 2016

Abstract. Privacy has been a major motivation for distributed problem optimization. However, even though several methods have been proposed to evaluate it, none of them is widely used. The Distributed Constraint Optimization Problem (DCOP) is a fundamental model used to approach various families of distributed problems. As privacy loss does not occur when a solution is accepted, but when it is proposed, privacy requirements cannot be interpreted as a criteria of the objective function of the DCOP. Here we approach the problem by letting both the optimized costs found in DCOPs and the privacy requirements guide the agents' exploration of the search space. We introduce Utilitarian Distributed Constraint Optimization Problem (UDCOP) where the costs and the privacy requirements are used as parameters to a heuristic modifying the search process. Common stochastic algorithms for decentralized constraint optimization problems are evaluated here according to how well they preserve privacy. Further, we propose some extensions where these solvers modify their search process to take into account their privacy requirements, succeeding in significantly reducing their privacy loss without significant degradation of the solution quality.

1 Introduction

In Distributed Constraint Optimization Problems (DCOP), agents have to find values to a set of shared variables while optimizing a cost function. To find such assignments, agents exchange messages (frequently assumed to have unspecified privacy implications) to explore the search space until an optimal solution is found or a termination condition is met. Thus, commonly agents reveal information during the solution search process, causing privacy to be a major concern in DCOPs [31].

The artificial intelligence assumption is that utility-based agents are able to associate each state with a utility value [24]. As such, the utility of each action is given by the difference between the utilities of final and initial states. If a user is concerned about privacy, then such a user can associate a utility value with the privacy of each piece of information in the definition of his local problem. If a user is interested in solving the problem, he must be also able to quantify the utility he draws from finding the solution. In a maximization DCOP we assume that the utility a user obtains from an assignment is represented by the values of the local constraints of the user for that assignment. Alternatively, with a minimization DCOP, the constraints would represent the costs. Certainly, these utilities can be modeled as a component in a multi-criteria DCOP [2].

Here we approach the problem by assuming that privacy has a utility that can be aggregated with the utility value for a given DCOP solution. We evaluate how much privacy is lost by the agents during the problem solving process, by the total utility of each information that was revealed. For DCOPs with private constraints one assumes that the cost/utility a constraint associate with a solution, is the kind of information that the agents would like to keep private. For DCOPs with privacy of domains, the existence of each value in the domain of a variable, would be kept private. For example, proposing an assignment with that value assigned to the variable has a privacy cost quantifying the desire of the agent to maintain its existence private. While sometimes possibilistic reasoning was used to guide search [27], in traditional algorithms agents explore the search space by proposing values as guided only by DCOP constraint costs. We propose a new DCOP framework with utility-based agents, where the utility of privacy as well as the utility of each solution is explicitly expressed. The framework is called Utilitarian Distributed Constraint Optimization Problem (UDCOP). Simple extensions to standard stochastic algorithms are studied to verify the impact of this interpretation of privacy.

Here we evaluate and compare several stochastic algorithms according to how well they preserve privacy. To do so, we generate distributed meeting scheduling (DMS) problems, as described in [18, 7]. In these problems, each agent own one variable, corresponding to the meeting to schedule. There exists a global constraint that requires all the variables to be equal, and also a unary constraint for each agent.

In the next section we discuss existing solvers and approaches to privacy for DCOPs. Further we formally define the concepts involved in UDCOPs. In Section 4 we introduce some extensions to common stochastic DCOP solvers that modify the search process to preserve privacy. We present our experimental results in Section 6, before presenting our conclusions.

2 Background

Let us first review the most relevant literature concerning DCOPs, stochastic algorithms and privacy measures.

2.1 Distributed Constraint Optimization Problems

Distributed Constraint Optimization Problems (DCOPs) have been extensively studied as a fundamental way of modeling combinatorial optimization problems in multi-agent systems. These problems have been addressed with a variety of algorithms, both stochastic and systematic. The systematic techniques range from highly asynchronous

[1] LAMIH UMR CNRS 8201, University of Valenciennes, France
[2] Nagoya Institute of Technology, Japan
[3] Kobe University, Japan
[4] Kyushu University, Japan
[5] Florida Institute of Technology, USA

protocols like ADOPT [21] or asynchronous branch and bound [30] to careful constraint pseudo-tree traversals like DPOP [23] or cluster exploitation like Asynchronous Partial Overlay [19]. Algorithms like ADOPT are known for their elegant treatment of searching within limited bounds from optima, while algorithms like DPOP are known for efficiently exploiting certain problem structures. The branch and bound algorithm [30] keeps expanding nodes in the search tree until a solution is found. For efficient use of memory, it keeps only the branch from the root node to the currently expanded node.

Another common algorithm is Synchronous Branch and Bound (SyncBB) [13], which was one of the first distributed algorithms for solving DCOPs. SyncBB organizes agents in a chain, and messages can traverse this chain upstream, or downstream. Some variations on the basic SyncBB algorithm include NCBB and AFB [8, 12].

DCOP problems have been addressed with constructive search [21, 3], fully cryptographic protocols [25], or hybrid crypto-constructive approaches for privacy [26, 9, 12]. Researchers have also addressed the issue of objective functions based on multiple criteria [4], as well as the impact of various aggregation functions for cost, ranging from the social welfare maximization of the pure addition to egalitarian leximin [22, 20].

2.2 Stochastic Algorithms

The main stochastic algorithms for solving DCOPs in practice are the distributed stochastic algorithm, the distributed simulated annealing, D-Gibbs, and the distributed breakout. In these algorithms, a flawed solution violating some constraints is revised until all constraints are satisfied.

Distributed Breakout The distributed breakout (DBO) [32] is an iterative improvement algorithm, originally proposed for DCOPs for hard constraints (distributed constraint satisfaction problems). In DBO, a weight starting at 1 is defined for each pair of assignments that does not satisfy some constraints. The evaluation of a given solution is the summation of the weights of all constraints for the involved assignment. With hard constraints, the summation is equal to the number of the constraint violations. In the breakout algorithm, an assignment is changed to decrease the solution value.

If the evaluation of the solution cannot be decreased by changing the value of any variable, the current state may be a local minimum. When trapped in a local minimum, the breakout algorithm increases the weights of constraint violation pairs in the current state by 1 so that the evaluation of the current state becomes higher than the neighboring states. Thus the algorithm can escape from a local minimum. Although the breakout algorithm is very simple, it is shown that it outperforms other iterative improvement algorithms.

Distributed Stochastic Algorithms The Distributed Stochastic Algorithm (DSA) is a family of algorithms [34]. In DSA, agents start by randomly selecting an initial value before entering a loop. In this loop, each agent first sends its new assigned value (if changed) to its neighbors, then it collects any new values assigned by those neighbors. Agents select the next candidate value based on the values received from other agents, and usually, based also on maximizing some utility function. The DSA family forms a baseline for evaluating other algorithms, and there exist a number of variations [34] of the DSA algorithm with slightly different properties. These variations differ mainly in the way they choose whether to keep the current state (assignment), or to assign a new one.

Stochastic algorithms are incomplete, namely not guaranteeing optimality. Other stochastic algorithms are Distributed Simulated Annealing and D-Gibbs. Distributed Simulated Annealing [1] differs from DSA in the way it picks the next value, and in the use of a schedule of temperatures to select the probabilities of changes to sub-optimal values. D-Gibbs [5] works by mapping DCOPs to probabilistic models and applies Markov Chain Monte Carlo.

2.3 Privacy

Privacy is a fundamental aspect in DCOPs, intrinsic to the main motivation, in addition to the usual efficiency/optimality trade-offs. The cost of privacy lost in the process of reaching a solution needs to be considered [10]. For example, in air traffic control [15], each airport has to allocate take-off and landing slots to the different flights. Such coordinated decisions are in conflict with the need to keep constraints private [6].

In existing works, several approaches have been developed to deal with privacy in DCOPs. The first approach using cryptographic techniques is [33]. While ensuring privacy [14], cryptographic techniques are usually slower, and sometimes require the use of external servers or computationally intensive secure function evaluation techniques that may not always be available or justifiable for their benefits [10].

Another family of approaches is based on using different search strategies to minimize privacy loss, as defined by certain privacy metrics.

Privacy categorization Agents might consider some -or all- of the following [11] as private information (that they rather not reveal), and a particular cost could incur in case any of them is revealed. Types of private information in DCOPs are: **domain privacy**, **constraint privacy**, **assignment privacy**, and **algorithmic privacy**.

A previously defined framework for modeling privacy requirements with DCOPs is the Valuations of Possible States (VPS). VPS [17, 16, 10] measures privacy loss by the extent to which the possible states of other agents are reduced [29]. Privacy is interpreted as a valuation on the other agents' estimates about the possible states that one lives in. During the search process, agents propose their values in an order of decreasing preference. At the end of the search process, the difference between the presupposed order of preferences and the real one observed during search determines the privacy loss: the greater the difference, the more privacy has been lost.

3 Concepts

In this section we define formally the distributed constraint optimization problem, as well as its extensions to utility-based agents.

3.1 Existing Frameworks

Let us start by presenting the DCOP framework and existing variations.

Distributed Constraint Optimization Problems The Distributed Constraint Optimization Problem (DCOP) is the formalism commonly used to model combinatorial problems distributed between several agents.

Definition 1. *A DCOP is a quadruplet* $\langle A, V, D, C \rangle$ *where:*

- $A = \langle A_1, ..., A_n \rangle$ *is a vector of* n *agents*

- $V = \langle x_1, ..., x_n \rangle$ *is a vector of* n *variables. Each agent* A_i *controls the variable* x_i*.*
- $D = \langle D_1, ..., D_n \rangle$ *is a vector of domains where* D_i *is the domain for the variable* x_i*, known only to* A_i*, and a subset of* $\{1, ..., d\}$*.*
- $C = \langle c_1, ..., c_m \rangle$ *is a vector of weighted constraints, each one defining a cost for each tuple of a relation between variables in* V*.*

The objective is to find an assignment for each variable that minimizes the total cost.

Example 1. *Suppose a problem concerning scheduling a meeting between three students. They all consider to agree on a place to meet on a given time, to choose between London, Madrid and Rome. For simplicity, in the next sections, we will refer to these possible values by their identifiers: 1, 2 and 3. The Student* A_1 *lives in Paris, and it will cost him* $\$70$*,* $\$230$ *and* $\$270$ *to attend the meeting in London, Madrid and Rome respectively. The Student* A_2 *lives in Berlin, and it will cost him* $\$120$*,* $\$400$ *and* $\$190$ *to attend the meeting in London, Madrid and Rome respectively. The Student* A_3 *lives in Brussels, and it will cost him* $\$40$*,* $\$280$ *and* $\$230$ *to attend the meeting in London, Madrid and Rome respectively. The objective is to find the meeting location that minimizes the total cost students have to pay in order to attend.*

The privacy costs for revealing her cost for locations 1, 2, and 3 for Student A_1 *are* $\$80, \$20, \$40$*. The privacy cost for locations 1, 2 and 3 are* $\$100, \$30, \$10$ *for Student* A_2 *and* $\$80, \$30, \$10$ *for Student* A_3*. There exist various reasons for privacy. For example, students may want to keep their cost for each location private, since it can be used to infer their initial location, and they would pay an additional (privacy) price rather than revealing the said travel cost. For example, Student* A_1 *associates* $\$50$ *privacy cost to the revelation of the travel cost of* $\$70$ *for meeting in London.*

DCOP The DCOP framework models this problem with:

- $A = \{A_1, A_2\, A_3\}$
- $V = \{x_1, x_2, x_3\}$
- $D = \{\{1, 2, 3\}, \{1, 2, 3\}\}$
- $C = \{\{(x_1 = 1), 70\}, \{(x_1 = 2), 230\}, \{(x_1 = 3), 270\},$
  $\{(x_2 = 1), 120\}, \{(x_2 = 2), 400\}, \{(x_2 = 3), 190\},$
  $\{(x_3 = 1), 40\}, \{(x_3 = 2), 280\}, \{(x_3 = 3), 230\}\}$
  $\{\neg(x_1 = x_2 = x_3), \infty\}\ \}$

where each constraint is described with the notation $\{p, c\}$ stating that if the predicate p holds then the cost c is payed, and the notation $(x = a)$ is a predicate stating that a variable x is assigned a value a.

With a DCOP, costs are paid when a solution is accepted. However, privacy costs are already paid whenever the corresponding assignments is proposed. This means that privacy costs cannot be interpreted as a criteria of a multi-objective DCOP. With a standard DCOP, agents can explore the search space, and then choose the solution with minimal cost. With privacy requirements, the exploration itself is costly, as it implies privacy leaks. This means that a given solution may imply different privacy loss depending on the algorithm used to reach the said solution. As it can be observed, DCOPs cannot model the details regarding privacy considerations.

One could attempt to model the privacy requirements by aggregating the solution quality, called $solutionCost$ and the $privacyCosts$ into a unique cost. However, this is not possible: In a DCOP, agents explore the search space to find a better solution, and only pay the corresponding solution cost when the search is over and the solution is accepted. This means that the solution cost decreases with time. However, privacy costs are cumulative and are paid during the search process itself (each time a solution is proposed), no matter what solution is accepted at the end of the computation. This means that the total privacy loss increases with time (see Figure 1). Aggregating the solution costs and privacy costs or using a multi-criteria DCOP would not consider the privacy cost of the solutions that are proposed but not kept as the final one.

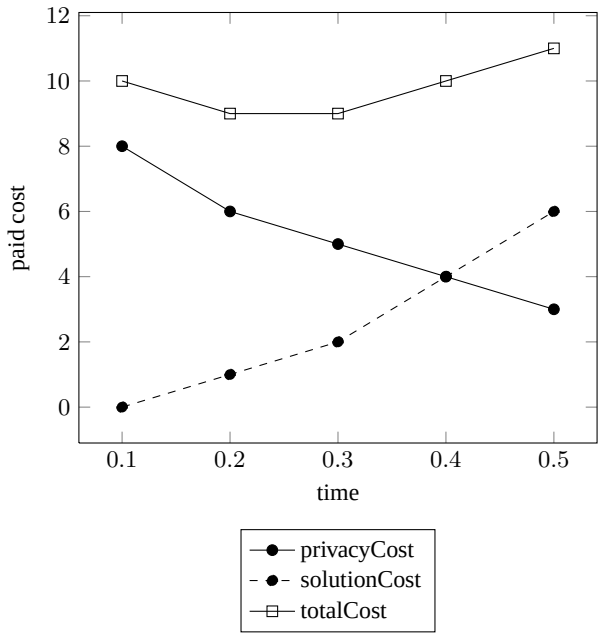

Figure 1. Evaluation of privacy loss on instances with different parameters.

3.2 Extensions

Utilitarian Distributed Constraint Optimization Problem We propose to ground the theory of DCOP in the well-principled theory of utility-based agentry. We introduce the Utilitarian Distributed Constraint Optimization Problem (UDCOP). Unlike previous DCOP frameworks, besides results, we are also interested in the search process.

Definition 2. *A UDCOP is a tuple* $\langle A, V, D, C, U \rangle$ *where:*

- $A = \langle A_1, ..., A_n \rangle$ *is a vector of* n *agents*
- $V = \langle x_1, ..., x_n \rangle$ *is a vector of* n *variables. Each agent* A_i *controls the variable* x_i*.*
- $D = \langle D_1, ..., D_n \rangle$ *is a vector of domains where* D_i *is the domain for the variable* x_i*, known only to* A_i*, and a subset of* $\{1, ..., d\}$*.*
- $C = \langle c_1, ..., c_m \rangle$ *is a vector of weighted constraints, each one defining a cost for each tuple of a relation between variables in* V*.*
- U*: a vector of privacy costs for each agent, each one defining the set of costs an agent suffers for the revelation of the values in his variable.*

The state *of agent* A_i *includes the subset of* D_i *that it has revealed, as well as the cost of the corresponding. The problem is to search for an assignment of the variables such such that the total utility is maximized (including privacy and solution utility/cost).*

Example 2. *The DCOP in the Example 1 is extended to a UDCOP by specifying the additional parameter* U*:*
$U = \langle \{u_{1,1} = 80, u_{1,2} = 20, u_{1,3} = 40\},$

$\{u_{2,1} = 100, u_{2,2} = 30, u_{2,3} = 10\}$,
$\{u_{3,1} = 80, u_{3,2} = 30, u_{3,3} = 10\}\rangle$

where $u_{i,j}$ is the privacy cost A_1 suffers from revealing the assignment $(x_i = j)$.

Note that to model this problem with the VPS framework, the 3 participants have to suppose an order of preference between all different possible values for each other agent. As agents initially do not know anything about others agents but the variable they share a constraint with, they have to suppose an equal distribution of all possible values for all other agents, meaning that they do not expect the feasibility of any value to be less secret, and so proposed first. In this direction one needs to extend VPS to be able to also model the kind of privacy introduced in this example.

UDCOPs with Private Constraints If agents are self-interested, each expecting a separate reward from the solution of the UDCOP, each of them potentially suffering personal costs described by constraints, and each of them having private costs for various configuration elements, a further extension would be needed.

Definition 3. *A UDCOP with Private Constraints (UDCOPPC) is a tuple $\langle A, V, D, C, U\rangle$ where:*

- $A = \langle A_1, ..., A_n\rangle$ *is a vector of n agents*
- $V = \langle x_1, ..., x_n\rangle$ *is a vector of n variables.*
- $D = \langle D_1, ..., D_n\rangle$ *is a vector of domains where D_i is the domain for the variable x_i.*
- $C = \langle C_1, ..., C_n\rangle$ *is a set of agent constraints, where each $C_i = \{c_{i,1}, ..., c_{i,m_i}\}$ is the set of weighted constraints known to agent A_i, each one defining a cost or utility for each tuple of a relation between variables in V.*
- $U = \langle U_1, ..., U_n\rangle$*: a vector of privacy costs for agents, each one defining the cost an agent suffers for the revelation of the weight it associates with a tuple in some of its constraint.*

The state *of agent A_i includes the subset of D_i that it has revealed, as well as the cost of the corresponding. The problem is to define a set of communication actions for each agent such that the total utility is maximized.*

Example 3. *The DCOP in the Example 1 is extended to a UDCOPPC by modifying the parameters C and U as follows:*

$C_1 = \{c_{1,1} = [(x_1 = 1), 70], c_{1,2} = [(x_1 = 2), 230],$
$c_{1,3} = [(x_1 = 3), 270], c_{1,4} = [\neg(x_1 = x_2 = x_3), \infty]\}$
$C_2 = \{c_{2,1} = [(x_2 = 1), 120], c_{2,2} = [(x_2 = 2), 400],$
$c_{2,3} = [(x_2 = 3), 190], c_{2,4} = [\neg(x_1 = x_2 = x_3), \infty]\}$
$C_3 = \{c_{3,1} = [(x_3 = 1), 40], c_{3,2} = [(x_3 = 2), 280],$
$c_{3,3} = [(x_3 = 3), 230], c_{3,4} = [\neg(x_1 = x_2 = x_3), \infty]\}$

$U_1 = \{(c_{1,1}, 80), (c_{1,2}, 20), (c_{1,3}, 40)\}$
$U_2 = \{(c_{2,1}, 100), (c_{2,2}, 30), (c_{2,3}, 10)\}$
$U_3 = \{(c_{3,1}, 80), (c_{3,2}, 30), (c_{3,3}, 10)\}$

A pair (c, v) appearing in the definition of the parameter U specifies that the privacy loss associated with the revelation of the cost/utility in constraint c is given by v.

The fact that a participant expects a reward r for finding a schedule for the meeting can be modeled in UDCOPPC by replacing the cost of the conflict in the constraint $[\neg(x_1 = x_2 = x_3), \infty]$ from infinity to r, obtaining $[\neg(x_1 = x_2 = x_3), r]$.

4 Algorithms

Now we discuss how the basic DBO and DSA algorithms are adjusted to UDCOPs. The state of an agent includes the `agent view`. After each state change, each agent computes the estimated utility of the state reached by each possible action, and selects randomly one of the actions leading to the state with the maximum expected utility.

In our algorithms, an information used by agents in their estimation of expected utilities is the risk of one of their assignments not being part of the final solution. For each agent A_i can be apriori estimated with the Equation 1:

$$futilityRisk = 1 - \frac{1}{|D_i|} \quad (1)$$

Before proposing a new value, agents estimate the utility that will be reached in the next state. This value is the summation of the costs of revealed `agent view`s (weighted by their probability to be the final solution) in the said state, and of the corresponding privacy costs.

If this $estimatedCost$ is lower than the estimation of the current state, the agent proposes the next value, otherwise it keeps its actual value.

The Distributed Breakout with Utility (DBOU) algorithm is obtained from DBO by adding the lines 2 to 8 in Algorithm 1. At line 2, the maximal improvement is initialized at 0. At line 3, the next value is initialized at the current value. At line 4, the possible next value is set to the value that gives the maximal improvement. At line 5, the set of revealed values is the union of the already revealed values and the new value At line 6, we estimate the cost reached after the next value is proposed. At line 7, the cost of the current state is estimated. At line 8, if the next cost is lower than the current cost, the maximal improvement and next value are updated

Similarly an algorithm called Distributed Stochastic Algorithm with Utilities (DSAU) is obtained from DSA, by adding the the lines 6 to 10 in Algorithm 2.

Example 4. *Continuing with Example 2, at the beginning of the computation with the DSAU solver, the participants select a random value. The resulting* `agent view` *of each agent is $x_1 = 1, x_2 = 1, x_3 = 3$. The utilities of the reached state are $70 + u_{1,1} = 70 + 80 = 150$, $120 + u_{2,1} = 120 + 100 = 220$, and $230 + u_{3,3} = 230 + 10 = 240$ for Student A_1, Student A_2, and Student A_3 respectively. The participants then inform each others of their value. They then consider changing their value to a new randomly selected one. The considered* `agent view` *is $x_1 = 2, x_2 = 3, x_3 = 1$. If the participants change their value, the utilities of the reached states would be $(70 + 230)/2 + u_{1,1} + u_{1,2} = 150 + 80 + 20 = 250$, $(120 + 190)/2 + u_{2,1} + u_{2,3} = 155 + 100 + 10 = 265$, and $(40 + 230)/2 + u_{3,3} + u_{3,1} = 135 + 10 + 80 = 225$, for Student A_1, Student A_2, and Student A_3 respectively. Student A_1 and Student A_2 do not propose the new value as it would increase their utility.*

However, Student A_3 chooses to change its value from 2 to 1 which lowers its utility from 240 to 225. In the next step, the `agent view` *is $x_1 = 1, x_2 = 1, x_3 = 1$. Participants then do not change their value anymore, as all other options would not decrease the utility. At the final step, the previous* `agent view` *is therefore the optimal solution. With DSAU, the reached utilities are $70 + 80 = 150, 120 + 100 = 220, 40 + 10 + 80 = 130$ for Student A_1, Student A_2, and Student A_3 respectively. With standard DSA, the final utilities are $(70 + u_{1,1} + u_{1,2} + u_{1,3} = 230$,*

$(120 + u_{2,1} + u_{2,2} + u_{2,3} = 260,$ *and* $(40 + u_{3,2} + u_{3,1} + u_{3,3} = 160,$ *for Student* A_1*, Student* A_2*, and Student* A_3 *respectively. Therefore, using DSAU instead of DSA reduces the utility by* $80, 40, 30$.

Algorithm 1: Procedure sendImprove in DBOU

Input:
Output:
currentEval = evaluation value of currentValue;
myImprove = 0 ;
newValue = currentValue ;
possibleValue = the value that gives the maximal improvement;
possibleRevealedConstraints = revealedConstraints + constraints containing possibleValue ;
nextCost = estimateCost($utilities$, $domain$, $nextRevealedValues$) ;
currentCost = estimateCost($utilities$, $domain$, $revealedValues$) ;
if *(nextCost* $<$ *currentCost)* **then**
    myImprove = possible max improvement ;
    newValue = the value that gives the maximal improvement;
if *currentEval = 0* **then**
    consistent = true
else
    consistent = false ;
    myTerminationCounter = 0;
if *myImprove* $>$ *0* **then**
    canMove = true;
    quasiLocalMinimum = false;
else
    canMove = false;
    quasiLocalMinimum = true;
send (improve, x_i, myImprove, currentEval, myTerminationCounter) to neighbors ;

Algorithm 2: DSAU algorithm

Input:
Output:
Randomly choose a value;
while *no termination condition is met* **do**
    if *a new value is assigned* **then**
        send the value to neighbors ;
    collect neighbors' new values, if any;
    possibleValue = randomly choose a value ;
    possibleRevealedConstraints = revealedConstraints + constraints containing possibleValue ;
    nextCost = estimateCost($utilities$, $domain$, $nextRevealedValues$) ;
    currentCost = estimateCost($utilities$, $domain$, $revealedValues$) ;
    if *(nextCost* $<$ *currentCost)* **then**
        assign possibleValue;

To adapt DBOU and DSAU for privacy of constraints in UDCOPPC, the revealed domains and possible revealed domains are changed to the revealed constraints and possible revealed constraints, respectively.

Algorithm 3: estimateCost

Input: $utilities$, $domain$, $revealedValues$
Output: $estimatedCost$
cost = 0 ;
privacyCost = 0 ;
foreach *value v in* $domain$ **do**
    foreach *constraint c in* $constraints$ **do**
        if *((c contains the assignment of v to* x_i*) and (v is in* $revealedValues$*))* **then**
            cost += (utilities.getCost(c) / (domain size of x_i) ;
            privacyCost += privacyCost of c ;
estimatedCost = cost + privacyCost ;
return estimatedCost ;

5 Discussion

To further clarify why Multi-Objective DCOPs (MO-DCOPs) cannot integrate our concept of privacy as one of the criteria they aggregate, we give an example of what would be achieved with MO-DCOPs, as contrasted with the results using the proposed UDCOPs.

Note that a MO-DCOP is a DCOP where the weight of each constraint tuple is a vector of values $[w_i]$, each value w_i representing a different metric. Two weights $[w_i^1]$ and $[w_i^2]$ for the same partial solution, inferred from disjoint sets of weighted constraints, are combined into a new vector $[w_i^3]$ where each value is obtained by summing the values in the corresponding position in the two input vectors, namely $w_i^3 = w_i^1 + w_i^2$. The quality of a solution of the MO-DCOP is a vector integrating the cost of all weighted constraints. The vectors can be compared using various criteria, such as leximin, maximin, social welfare or the Theil index [22, 20].

In the following example we show a comparative trace based on one of the potential techniques in MO-DCOPs, to provide a hint on why MO-DCOPs cannot aggregate privacy lost during execution in the same way as UDCOP. In this example, the privacy value of each assignment and its constraint cost are two elements of an ordered pair defining the weight of the MO-DCOP. For illustration, in this example pairs of weights are compared lexicographically with the privacy having priority.

Example 5. *Suppose we now want to model the Example 2 with a MO-DCOP. As also illustrated in the trace in Table 1, at the beginning of the computation with the DSA solver, the participants select a random value. The resulting* `agent view` *is* $x_1 = 1, x_2 = 1, x_3 = 3$*. The participants then inform each others of their value. They then consider changing their value to a new randomly selected one. The considered* `agent view` *is* $x_1 = 2, x_2 = 3, x_3 = 1$.

Like with UDCOPs, Student A_1 *does not propose the new value as it would increase their cost, and Student* A_3 *chooses to change its variable's value from* 2 *to* 1.

However, with MO-DCOPs Student A_2 *changes its value to* 3*, which is not the case with UDCOPs, which implies privacy loss. The* `agent view` *is now* $x_1 = 1, x_2 = 3, x_3 = 1$.

As we see, with the MO-DCOP model, Student A_2 *reveals more values and loses more privacy (with 110-100=10 units of privacy more) than with UDCOPs.*

Framework	UDCOP			MO-DCOP		
Agent	Student A_1	Student A_2	Student A_3	Student A_1	Student A_2	Student A_3
value step 1	1	1	3	1	1	3
cost	70	120	230	70	120	230
privacyCost	80	100	10	80	100	10
situation	150	220	240	$[80, 70]$	$[100, 120]$	$[10, 230]$
believed next state						
considered	2	3	1	2	3	1
cost	150	155	135	230	190	40
privacyCost	100	110	90	20	10	80
situation	250	265	225^*	$[20, 230]^*$	$[10, 190]^*$	$[80, 40]$
achieved next state						
value step 2	1	1	1	2	3	3
cost	70	120	40	230	190	230
privacyCost	80	100	90	100	110	10
situation	150	220	130	$[100, 230]$	$[110, 190]$	$[10, 230]$

Table 1. Comparative trace of two rounds with UDCOP DSAU vs. MO-DCOP DSA with lexicographical comparison on vectors, privacy first. Candidate values are marked with * if they are better than old values, and will be adopted.

6 Experimental Results

We evaluate our framework and algorithms on randomly generated instances of *distributed meeting scheduling problems (DMS)*. Previous work [28] in distributed constraint satisfaction problems has already addressed the question of privacy in distributed meeting scheduling by considering the information on whether an agent can attend a meeting to be private. They evaluate the privacy loss brought by an action as the difference between the cardinalities of the final set and of the initial set of possible availabilities for a participant.

The algorithm we use to generate the problem is:

1. We create the variables (one per *participant agents*).
2. We initialize their domain (possible *times*).
3. We add the global constraint *"all equals"*.
4. Unary constraints are added to variables, to fit the density.
5. For each unary constraint, we generate a cost between 0 and 9.
6. For each value of each variable, we generate a revelation cost uniformly distributed between 0 and 9.

The experiments are carried out on a computer under Windows 7, using a 1 core 2.16GHz CPU and 4 GByte of RAM. In Figure 2, we show the total amount of privacy lost by all agents, averaged over 50 problems, function of the density of unary constraints. In Figure 3, we show the total cost (solutionCost + privacyCost) for all agents. Table 2 shows that while DBOU is leading to slightly higher solution cost, the solution cost for DSAU is the same as for DSA. As we can see, for all problems, the curves of DBOU and DSAU are similar to the ones of DBO and DSA, respectively, meaning that adding privacy preservation techniques with the UDCOP framework does not degrade the quality of the solution. The problems are parametrized as follows: 10 agents, 10 possible values, the cost for the constraints is a random number between 0 and 9, and the cost of a revelation is a random number between 0 and 9. Each set of experiments is an average estimation over 50 instances with the different algorithms (i.e, DBO, DBOU, DSA, DSAU).

Algorithm	DBO	DBOU	DSA	DSAU
Average Solution Quality	4.33	4.55	5.11	5.11

Table 2. Average solution quality per agent for various algorithms.

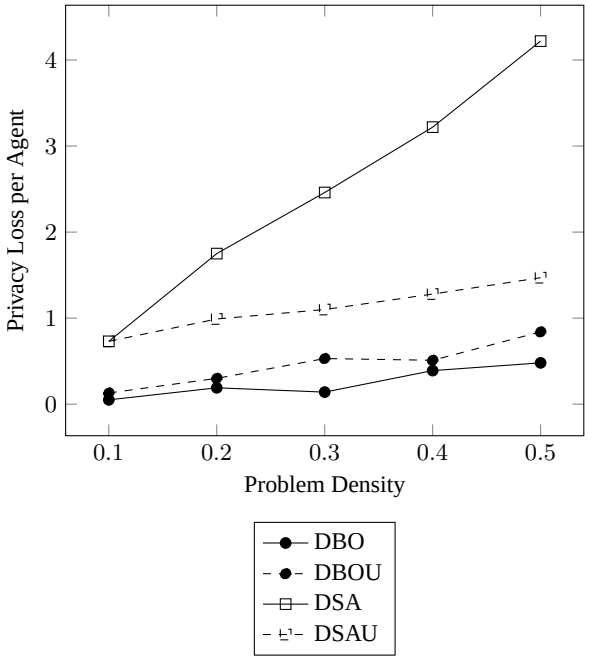

Figure 2. Evaluation of privacy loss on instances with different densities and algorithms.

7 Conclusion

While various previous efforts have addressed privacy in distributed constraint optimization problems, none of the existing techniques is widely used, likely due to the difficulty in modeling common problems. As privacy cannot be interpreted as a criteria of a standard DCOP, we propose in this article a framework called Utilitarian Distributed Constraint Optimization Problem (UDCOP). It models the privacy loss for the revelation of an agent's costs for violating constraints. We present algorithms that let agents use information about privacy to modify their behavior and guide their search process, by proposing values that reduce the amount of privacy loss. We then show how adapted stochastic algorithms (DBOU and DSAU) behave and compare them with standard techniques on different types of distributed meeting scheduling problems. The experiments show that

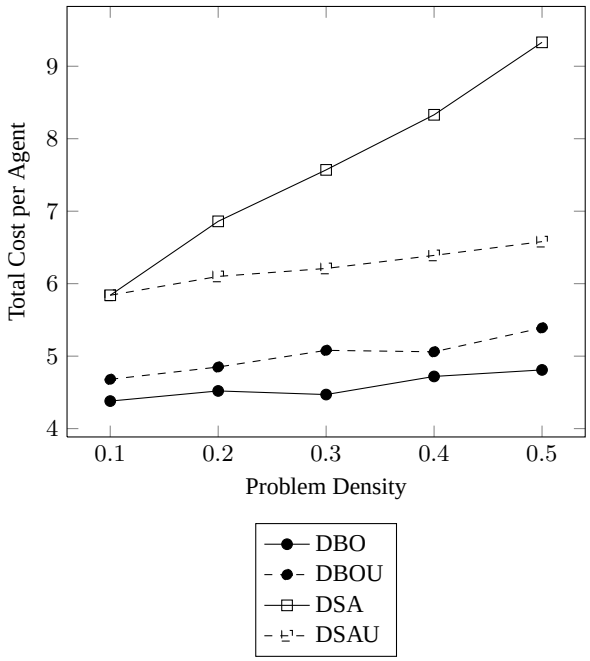

Figure 3. Evaluation of total cost on instances with different densities and algorithms.

explicit modeling and reasoning with the utility of privacy allows for significant savings in privacy with minimal impact on the quality of the achieved solutions.

REFERENCES

[1] Muhammad Arshad and Marius C Silaghi, 'Distributed simulated annealing', *Distributed Constraint Problem Solving and Reasoning in Multi-Agent Systems*, **112**, (2004).

[2] Emma Bowring, Milind Tambe, and Makoto Yokoo, 'Distributed multi-criteria coordination: Privacy vs. efficiency', in *Nineteenth International Joint Conference on Artificial Intelligence (IJCAI-2005)-Workshop on Distributed Constraint Reasoning-DCR05, Edinburgh, Scotland*. Citeseer, (2005).

[3] Ismel Brito, Amnon Meisels, Pedro Meseguer, and Roie Zivan, 'Distributed constraint satisfaction with partially known constraints', *Constraints*, **14**(2), 199–234, (2009).

[4] Maxime Clement, Tenda Okimoto, Tony Ribeiro, and Katsumi Inoue, 'Model and algorithm for dynamic multi-objective distributed optimization', in *PRIMA 2013: Principles and Practice of Multi-Agent Systems*, 413–420, Springer, (2013).

[5] H. C. Lau D. T. Nguyen, W. Yeoh, 'Distributed gibbs: A memory-bounded sampling-based dcop algorithm', in *AAMAS*, pp. 167–174, (2013).

[6] Boi Faltings, Thomas Léauté, and Adrian Petcu, 'Privacy guarantees through distributed constraint satisfaction', in *Web Intelligence and Intelligent Agent Technology, 2008. WI-IAT'08. IEEE/WIC/ACM International Conference on*, volume 2, pp. 350–358. IEEE, (2008).

[7] A. Gershman, A. Grubshtein, A. Meisels, L. Rokach, and R. Zivan, 'Scheduling meetings by agents', in *Proc. 7th International Conference on Practice and Theory of Automated Timetabling (PATAT 2008). Montreal (August 2008)*, (2008).

[8] A. Gershman, A. Meisels, and R. Zivan, 'Asynchronous forward-bounding for distributed cops', *JAIR*, **34**, 61–88, (2009).

[9] Rachel Greenstadt, Barbara Grosz, and Michael D Smith, 'Ssdpop: improving the privacy of dcop with secret sharing', in *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, p. 171. ACM, (2007).

[10] Rachel Greenstadt, Jonathan P Pearce, and Milind Tambe, 'Analysis of privacy loss in distributed constraint optimization', in *AAAI*, volume 6, pp. 647–653, (2006).

[11] Tal Grinshpoun. When you say (dcop) privacy, what do you mean, 2012.

[12] Tal Grinshpoun and Tamir Tassa, 'A privacy-preserving algorithm for distributed constraint optimization', in *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pp. 909–916. International Foundation for Autonomous Agents and Multiagent Systems, (2014).

[13] Katsutoshi Hirayama and Makoto Yokoo, 'Distributed partial constraint satisfaction problem', in *Principles and Practice of Constraint Programming-CP97*, 222–236, Springer, (1997).

[14] Martin Hirt, Ueli Maurer, and Bartosz Przydatek, 'Efficient secure multi-party computation', in *Advances in Cryptology—ASIACRYPT 2000*, 143–161, Springer, (2000).

[15] International Air Transport Association IATA, *Worldwide scheduling guidelines*, International Air Transport Association, 2005.

[16] Rajiv T Maheswaran, Jonathan P Pearce, Emma Bowring, Pradeep Varakantham, and Milind Tambe, 'Privacy loss in distributed constraint reasoning: A quantitative framework for analysis and its applications', *Autonomous Agents and Multi-Agent Systems*, **13**(1), 27–60, (2006).

[17] Rajiv T Maheswaran, Jonathan P Pearce, Pradeep Varakantham, Emma Bowring, and Milind Tambe, 'Valuations of possible states (vps): a quantitative framework for analysis of privacy loss among collaborative personal assistant agents', in *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pp. 1030–1037. ACM, (2005).

[18] Rajiv T Maheswaran, Milind Tambe, Emma Bowring, Jonathan P Pearce, and Pradeep Varakantham, 'Taking dcop to the real world: Efficient complete solutions for distributed multi-event scheduling', in *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pp. 310–317. IEEE Computer Society, (2004).

[19] Roger Mailler and Victor R Lesser, 'Asynchronous partial overlay: A new algorithm for solving distributed constraint satisfaction problems', *Journal of Artificial Intelligence Research*, 529–576, (2006).

[20] Toshihiro Matsui, Marius Silaghi, Tenda Okimoto, Katsutoshi Hirayama, Makoto Yokoo, and Hiroshi Matsuo, 'Leximin asymmetric multiple objective dcop on factor graph', in *PRIMA 2015: Principles and Practice of Multi-Agent Systems*, 134–151, Springer, (2015).

[21] Pragnesh Jay Modi, Wei-Min Shen, Milind Tambe, and Makoto Yokoo, 'Adopt: Asynchronous distributed constraint optimization with quality guarantees', *Artificial Intelligence*, **161**(1), 149–180, (2005).

[22] A. Netzer and A. Meisels, 'Social dcop - social choice in distributed constraints opti- mization', in *International Symposium on Intelligent Distributed Computing*, pp. 35–47, (2011).

[23] Adrian Petcu and Boi Faltings, 'A scalable method for multiagent constraint optimization', Technical report, (2005).

[24] Stuart Russell and Peter Norvig, 'Artificial intelligence: a modern approach', (2010).

[25] Marius-Calin Silaghi, Boi Faltings, and Adrian Petcu, 'Secure multi-party constraint optimization simulating DFS tree-based variable elimination', in *ISAIM*, (2006).

[26] Tamir Tassa, Roie Zivan, and Tal Grinshpoun, 'Max-sum goes private', in *Proceedings of the 24th International Conference on Artificial Intelligence*, pp. 425–431. AAAI Press, (2015).

[27] R Wallace and Marius C Silaghi, 'Using privacy loss to guide decisions in distributed csp search', in *FLAIRS'04*, (2004).

[28] Richard J Wallace and Eugene C Freuder, 'Constraint-based reasoning and privacy/efficiency tradeoffs in multi-agent problem solving', *Artificial Intelligence*, **161**(1), 209–227, (2005).

[29] R.J. Wallace and E.C. Freuder, 'Constraint-based multi-agent meeting scheduling: Effects of agent heterogeneity on performance and privacy loss', in *DCR*, pp. 176–182, (2002).

[30] William Yeoh, Ariel Felner, and Sven Koenig, 'Bnb-adopt: An asynchronous branch-and-bound dcop algorithm', in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*, pp. 591–598. International Foundation for Autonomous Agents and Multiagent Systems, (2008).

[31] Makoto Yokoo, Edmund H Durfee, Toru Ishida, and Kazuhiro Kuwabara, 'The distributed constraint satisfaction problem: Formalization and algorithms', *Knowledge and Data Engineering, IEEE Transactions on*, **10**(5), 673–685, (1998).

[32] Makoto Yokoo and Katsutoshi Hirayama, 'Distributed breakout algorithm for solving distributed constraint satisfaction problems', in *Proceedings of the Second International Conference on Multi-Agent Sys-*

tems, pp. 401–408, (1996).

[33] Makoto Yokoo, Koutarou Suzuki, and Katsutoshi Hirayama, 'Secure distributed constraint satisfaction: Reaching agreement without revealing private information', in *Principles and Practice of Constraint Programming-CP 2002*, pp. 387–401. Springer, (2002).

[34] Weixiong Zhang, Guandong Wang, and Lars Wittenburg, 'Distributed stochastic search for constraint satisfaction and optimization: Parallelism, phase transitions and performance', in *Proceedings of AAAI Workshop on Probabilistic Approaches in Search*, (2002).